arXiv:1402.0096v1 [math.NA] 1 Feb 2014

Adapted Basis for Non-Local Reconstruction of Missing Spectrum

Antonin Chambolle*, Khalid Jalalzai**

CMAP, École Polytechnique, CNRS

Abstract. The object of this work is to design an adequate regularization for the problem of recovering missing Fourier coefficients, particularly in some non standard situations were low frequency coefficients are lost. In the framework of non-local regularization, we propose a technique to build an original patchwise similarity measure that is adapted to the missing spectrum. Then, a simple Non-Local quadratic energy is minimized. By construction, the similarity criterion is invariant under the corruption process so that the distance between two patches of the corrupted image is almost exactly equal to the one computed on the clean image. We illustrate our method with experiments which show its efficiency, both in terms of speed and quality of the results, with respect to other common approaches. We show that the method is practical on synthetic examples which are built upon models of inverse scattering problems, synthetic aperture mirrors for spatial imaging or also medical imaging.

1 Introduction

In this paper, we consider the problem of retrieving missing Fourier coefficients of a raw data. This problem has important applications in various fields:

- The zooming problem in image processing where one has to construct high frequency data from a low-resolution sample. This is of some importance nowadays for the transition of SD videos to HD ones.
- Aperture Synthesis for spatial imaging where the corruption process is given by a mask whose typical example is shown in Fig. 4. (For further details see [18].)
- The tomography problem for medical imaging or seismic imaging. In this case Fourier coefficients usually lie on straight lines that are either parallel or that cross at the origin.
- The inverse scattering problem where one is interested in recovering the shape of an hidden object using electromagnetic or acoustic waves.

Let us consider the latter example (see [7] for details). Let an *incident acoustic plane wave* $u^i(x) = e^{ikx\cdot d}$ propagate in the direction of the unit vector d in an

* antonin.chambolle@cmap.polytechnique.fr
** khalid.jalalzai@polytechnique.edu

isotropic medium. Here $k > 0$ is the *wave number*. In case there is an inhomogeneity D (the hidden object), the wave will be "scattered" and give rise to another wave u^s. The latter has an asymptotic behavior

$$u^s(x) = \frac{e^{ik|x|}}{|x|} u_\infty(\hat{x}, d) + O\left(\frac{1}{|x|^2}\right), \tag{1}$$

where $u_\infty(\hat{x}, d)$ is known as the *far field pattern* and gives an idea of the behavior of the scattered wave at large distance. Here, $\hat{x}$ is the observation direction and d is the incident wave direction. The *direct problem* amounts to find $u_\infty(\hat{x}, d)$ whereas the *nonlinear inverse problem* takes the direct method as a starting point and asks what is the nature of the scatterer D that gave rise to such a farfield. In what follows, we assume that k is sufficiently small so that the *Born approximation* implies that

$$u_\infty(\hat{x}, d) \approx \int_{\mathbb{R}^N} \chi_D(y) e^{-ik(\hat{x}-d)\cdot y} dy. \tag{2}$$

This means that $u_\infty(\hat{x}, d)$ can be interpreted as Fourier coefficients in the ball of radius $2k$. In practice, having only limited incident waves and measurements we end up with a sampling of the spectrum of $\chi_D(y)$ that is depicted in Fig. 1.

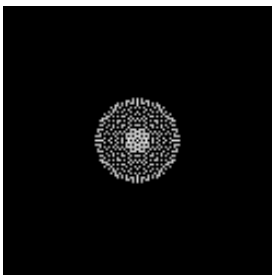

Fig. 1. Spectrum obtained with 32 incident planewaves and 32 measurement directions.

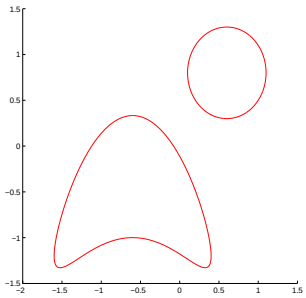

Fig. 2. Scatterer D.

In general non regular sampling of the spectrum is a difficult problem (see for instance [23]). Though if D is bounded, its Fourier transform is a C^∞ function that can be interpolated on a uniform grid. This allows to use the `fft`.

In [14, 24], the authors consider the Total Variation (henceforth denoted TV) to recover missing Fourier coefficients. Let us compare the performance of this approach with that of the factorization method [21] thanks to a numerical test. We consider the object D that is bounded by the red curves in Fig. 2. The inverse problem is then solved given 32 incident waves and 32 measurement directions. The wave number is $k = 3\pi$. The factorization method, which performs quite well for small values of k, yields quite poor results in this extreme case:

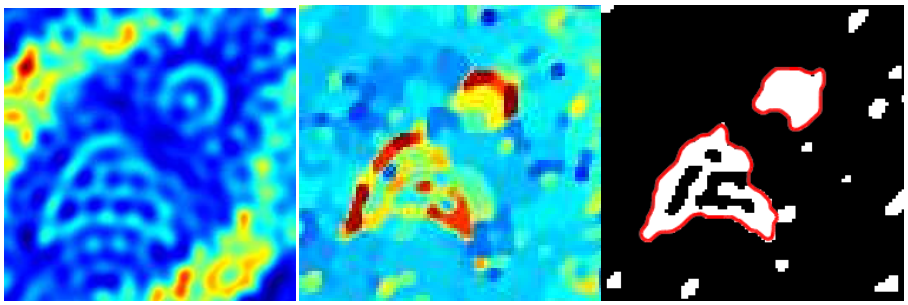

Fig. 3. From left to right, scatterer obtained by the *factorization method* [21], scatterer obtained by Total Variation minimization, the result is then thresholded and segmented [6] (red curve on the right).

The TV-based method yields really good results in this case. Indeed, from partial and noisy measurements, we were able to distinguish two objects that are merely separated by $3/4\lambda$ where $\lambda = 2\pi/k$ is the *wavelength*. This value is close to the theoretical diffraction limit that is $\lambda/2$. Note also it is not necessary to tune a Lagrange multiplier as it is usually done. Indeed, if one denotes M the discrete set of points where u_∞ is known, one can simply minimize the constrained problem

$$\min_{\mathcal{F}u_{|M}=u_\infty} TV(u), \tag{3}$$

to get the restoration shown in Fig. 3. Total Variation was already used for solving the nonlinear inverse scattering problem in [29]. The idea of considering the inverse problem within the Born approximation as an inpainting problem in the Fourier domain is discussed in [10].

This example shows that using variational methods for the general spectrum interpolation problem may yield quite satisfactory results. We should also mention the not so different problem of restoring missing wavelet coefficients which is treated in [5, 30]. For such problems, can we do better than using TV? Some recent papers address the problem of solving general inverse problems in imaging by means of non-local methods. These methods obviously apply to the problem of retrieving missing Fourier coefficients.

2 Prior Works

Traditional methods in image processing are based on local properties of images (Wavelets, Total Variation). Recently, state of the art results were achieved for the denoising problem by Non-Local methods that exploit redundancies in images. The idea of using self-similarities in images was first exploited in [3]. They proposed a filter that averages similar patterns of a noisy image $g = g_0 + n$

defined on $\Omega \subset \mathbb{R}^2$ even though these self-similarities occur at large distance. The filter reads as follows

$$\text{NLMeans}(g)(x) = \frac{1}{C(x)} \int_\Omega g(y) \exp\left(-\frac{\|p_g(x) - p_g(y)\|_2}{h}\right) dy, \tag{4}$$

where $p_g(x),\ p_g(y)$ are patches of g and h determines the selectivity of the similarity measure. Many modifications of this filter were considered for denoising purposes (adaptive h, adaptive window [19], shape adaptive patches [9]).

The use of patches has been widely adopted in the image processing community and in the recent years, these non-local methods were extended to the study of general inverse problems (see [20, 13, 12, 26, 1]). Such a general inverse problem can be modeled as follows $g = Ag_0 + n$. Here g_0 is the original image defined on $\Omega \subset \mathbb{R}^2$, A is a linear transformation and n is a white Gaussian noise. The classical local methods were adapted by considering a *pointwise Non-Local regularization* as introduced in [13]

$$J_w(u) = \int_{\Omega\times\Omega} |u(x) - u(y)|^\alpha w(x,y)\eta(x-y)dxdy \tag{5}$$

for some $\alpha \geq 1$ and where $w(x,y) = \exp\left(-\frac{\|p_g(x)-p_g(y)\|_2}{h}\right)$ is the weight function used in the NLMeans filter and is based on the Sum of Squared Differences (henceforth called SSD distance). As for η, it is a function of compact support centered at the origin. It indicates that patches that are too far from each other should not be taken into account. Indeed, it was observed that the Non-Local methods yield much better results if one seeks for similar patches in the support of η, referred to as the *search window* (see also assumption (ii) in Section 7).

Peyré et al. proposed in [27] (see also [1]) a *patchwise Non-Local regularization*

$$J_w(u) = \int_{\Omega\times\Omega} \|p_u(x) - p_u(y)\|_2 w(x,y)\eta(x-y)dxdy\,. \tag{6}$$

To get a restored image one then has to minimize the following Non-Local energy

$$\mathcal{E}(u) = \frac{1}{2}\|Au - g\|_2^2 + \lambda J_w(u). \tag{7}$$

It is further remarked that one can enhance the restoration by recomputing the weight w regularly. In [1, 27], the authors proposed a general framework where the weight $w(x,\cdot)$ is interpreted up to renormalization as a density of probability and is an unknown of the problem. Then the energy to be minimized involves the potential energy of $w(x,\cdot)$, used to infer unknown probability distributions. The resulting energy is not convex in w and an alternate coordinate descent gives back the SSD-based weight but computed this time on the u being processed. This actually corresponds to the aforementioned weight recomputation procedure.

3 Contribution

In this paper, we deal with the problem of restoring missing Fourier coefficients of an image. In particular we propose a technique to build a distance insensitive

to the degradation of an image. The corruption process is supposed to be known. We show that using this distance between patches gives very good reconstruction results by solving a simple quadratic energy.

In the classical Non-Local approaches, a critical step is to compute the SSD distance between two patches of the corrupted image. This strategy dates back to the NLMeans [3].We explain how to replace this step by the computation of a simpler ℓ^2-type distance that is really adapted to the problem of spectrum reconstruction and does not incorporate spurious information. The idea is to define atoms similar to Gabor filters to test whether two regions of the corrupted image are similar. These atoms should be as concentrated as possible and should not depend on the corruption process. This way two patches that are close in the clean image will be close in the corrupted image. We now first introduce our variational framework, which is relatively standard.

4 A Non-Local Energy for the Problem

Henceforth, we work in $\mathbb{R}^2$ which will is equipped with the norm $\max\{|x_1|, |x_2|\}$ so that "balls" are in fact squares.

In this section, we first formalize the problem of reconstructing unknown coefficients of a Fourier series. Let us call $M \subset \mathbb{Z}^2$ the finite mask of points where the Fourier coefficients are known. We shall assume that it is symmetric with respect to the origin. Assuming that the image is periodic and defined on $\mathbb{T} = [0,1]^2$, one thus considers respectively the clean image and the corrupted images

$$g_0(x) = \sum_{k \in \mathbb{Z}^2} c_k e^{-2i\pi k \cdot x}, \; g(x) = \sum_{k \in M} c_k e^{-2i\pi k \cdot x}. \tag{8}$$

We assume that not only the high frequencies but also middle range frequencies are lost in the corruption process. In other words, we keep the Fourier coefficients corresponding to the white areas of the mask M. The typical M we will consider here is the following

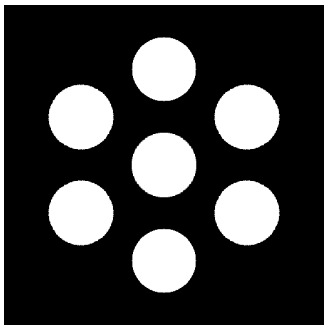

Fig. 4. Corruption mask M.

for which usual *local* TV-based reconstruction techniques totally fail at recovering the middle-range frequencies, see for instance Fig. 5, right.

In the sequel, we introduce the two subsets of L^2: $\mathcal{M} = \{\sum_{k\in M} c_k e^{-2i\pi k\cdot x}\}$ and its complement $\mathcal{M}^\perp = \{\sum_{k\notin M} c_k e^{-2i\pi k\cdot x}\}$ where in both cases $(c_k)_k$ runs over all sequences in $\ell^2(\mathbb{Z}^2;\mathbb{R})$, which are even in k. Let us denote $P_{\mathcal{M}}$ the orthogonal projection on $\mathcal{M}$. We recall that the Fourier transform on the torus $\mathcal{F} : L^2(\mathbb{T}) \to \ell^2(\mathbb{Z}^2)$ is an isometry and thus invertible so the corrupted data $g \in \mathcal{M}$ that one has to interpolate is obtained by

$$g = \mathcal{F}^{-1} \circ P_{\mathcal{M}} \circ \mathcal{F}(g_0) = \mathcal{F}^{-1}(\chi_M \mathcal{F}(g_0)), \tag{9}$$

We shall call *patch* centered at $x_k \in \mathbb{R}^2$, denoted $p_k(g)$, the image $g(\cdot - x_k)\varphi$ where φ is a test function with support $C(0, \frac{\rho}{2}) := \left[-\frac{\rho}{2}, \frac{\rho}{2}\right]^2$ and ρ is the *patch size*.

For the moment let us assume that the original image is such that there are two distinct $x_k, x_\ell \in \mathbb{R}^2$ with $g_0(x - x_k)\varphi(x) = g_0(x - x_\ell)\varphi(x),\ \forall x \in \mathbb{R}^2$. This is to say that two patches are similar in the original image. Then clearly, the part of the image v corresponding to the missing frequencies is among the solutions of

$$\min_{v\in\mathcal{M}^\perp} \int_{\mathbb{T}} \psi(x)^2 |(g+v)(x-x_k) - (g+v)(x-x_\ell)|^\alpha dx \tag{10}$$

where ψ is smooth and such that $\mathrm{supp}(\psi) \subset \mathrm{supp}(\varphi) = C(0, \frac{\rho}{2})$, and $\alpha \geq 1$ is fixed. This means that knowing that $p_k(g_0)$ and $p_\ell(g_0)$ are similar one can hope to get a restored spectrum by minimizing (10). The reconstruction is obviously not unique since modifying v out of $(x_k + \mathrm{supp}(\psi)) \cup (x_\ell + \mathrm{supp}(\psi))$ does not change the energy (and numerical simulations actually show that the reconstructed v has the same support as $(x_k + \mathrm{supp}(\psi)) \cup (x_\ell + \mathrm{supp}(\psi)))$. Thus, we have to take into account all the patches in the image to get a global reconstruction.

We are therefore led to consider the following problem

$$\min_{v\in\mathcal{M}^\perp} \int_{\mathbb{T}} \sum_{(k,\ell)\in I} \psi^2 |(g+v)(x-x_k) - (g+v)(x-x_\ell)|^\alpha w(x_k, x_\ell)\, dx \tag{11}$$

where $w(x_k, x_\ell) = \exp\left(-\frac{\delta(x_k,x_\ell)}{h}\right)$ and $\delta(x_k, x_l)$, that is going to be defined in the sequel, should tell us from the corrupted image g whether one had for the original image $p_{g_0}(x_k) \sim p_{g_0}(x_\ell)$.

There are two interesting values for α, namely $\alpha = 1$ and $\alpha = 2$. Our first experiment (Fig. 5) shows that the reconstruction with (11) for both these choices is almost perfect if the weights $w(x_k, x_\ell)$ are derived from the "oracle" distance $\delta(x_k, x_\ell) = \|p_{g_0}(x_k) - p_{g_0}(x_\ell)\|_2$ given by the SSD distance of the patches of the (normally unknown) original image g_0.

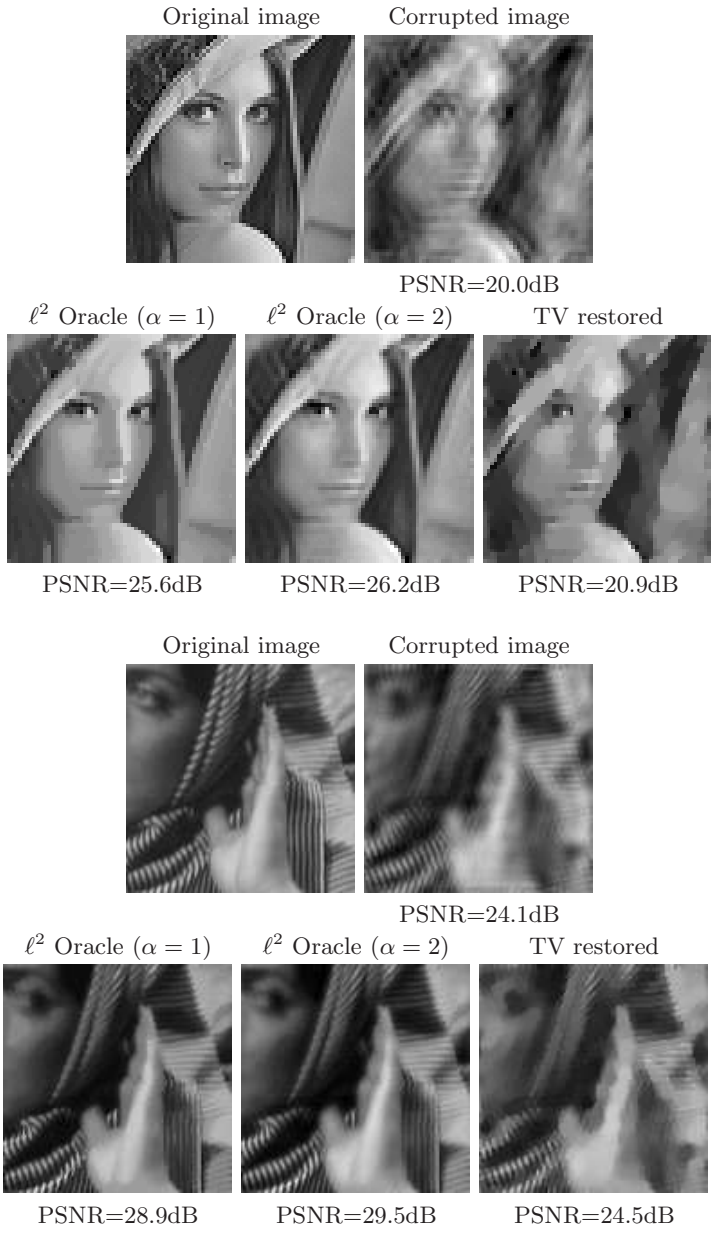

Fig. 5. Non-Local restoration thanks to the ideal distance. Both cases $\alpha = 1$ and $\alpha = 2$ are considered.

From these experiments, it seems clear that both values of α allowed us to produce a cleaner image. Due to the huge size of our problems, we will stick in general to the case of the exponent $\alpha = 2$ which *de facto* excludes the case of standard Non-Local Total Variation. We will see that this choice leads to

computationally tractable problems with excellent results, for some masks (such as the one in Fig. 4), provided the distance between similar patches is correctly estimated. For the considered mask (Fig. 4), standard Total Variation yield low quality reconstructions (right).

5 Construction of Atoms Adapted to the Corruption

Our aim is to define a similarity measure for patches that is not modified through the corruption process $\mathcal{F}^{-1} \circ P_{\mathcal{M}} \circ \mathcal{F}$. This way, if two points are close in the original image they will remain close in the perturbed image.

To do so, our idea is to find a family of test functions $(\phi_\beta)_\beta$ dense in $\mathcal{M}$, that do not depend on g and such that

$$g * \phi_\beta = g_0 * \phi_\beta, \ \forall \beta. \tag{12}$$

The atoms $(\phi_\beta)_\beta$ should obviously depend on the mask M. Indeed, for any $g \in \mathcal{M}$,

$$\langle g, \phi_\beta \rangle = \langle \mathcal{F}^{-1}(\chi_M \mathcal{F}(g)), \phi_\beta \rangle = \langle g, \mathcal{F}^{-1}(\chi_M \mathcal{F}(\phi_\beta)) \rangle \tag{13}$$

which means that $\mathcal{F}\phi_\beta = \chi_M \mathcal{F}\phi_\beta$ and therefore $\text{supp}(\mathcal{F}\phi_\beta) \subset M$ for any β. From the uncertainty principle (see [22]), one knows that ϕ_β is not of compact support (since M is) but one could still expect it to be localized in space. This can be ensured by defining these functions as minimizers of the following parameterized problems

$$\phi_\beta = \text{argmin} \left\{ \int_\Omega |\phi(x)|^2 |x|_2^p dx, \phi \in \mathcal{M}, \|\phi\|_2 = 1, \phi \perp \text{Span}\{\phi_{\beta'}, \beta' < \beta\} \right\} \tag{14}$$

where we minimize the p-moments of ϕ, for some $p > 1$.

We can then consider a measure of similarity that is of the form

$$\delta(x_k, x_l) = \left(\sum_{\beta < \beta_0} |g * \phi_\beta(x_k) - g * \phi_\beta(x_\ell)|^2 \right)^{\frac{1}{2}}, \tag{15}$$

where β_0 sets how localized the considered atoms ϕ_β are. Now, by definition, the similarity measure δ is invariant under the corruption process $\mathcal{F}^{-1} \circ P_{\mathcal{M}} \circ \mathcal{F}$. In other words, two patches that were identical in the original image g_0 may not necessarily match according to the SSD but will match according to the atom-based distance, provided that the atoms are well localized.

For the mask M of Fig. 4, the first atoms are actually really localized and can be used as test functions. Let us have a closer look at the 7 first atoms, computed for $p = 4$, with their respective spectra:

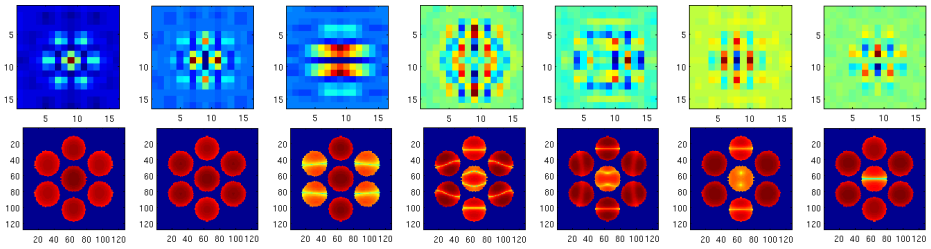

Fig. 6. First line: Atoms $(\phi_n)_{n=1,...,7}$ with $p = 4$ adapted to the 128×128 mask M zoomed in. Second line: Their respective spectra $\log(|\mathcal{F}(\phi_n)|)$, $n = 1, ..., 7$.

These atoms adapted to the mask M may recall the reader of the Gabor atoms (see [22]). However, they have the advantage of having a prescribed spectrum and being as localized as possible. They are in a sense the optimal functions satisfying these two conditions (see Section 7 for a discussion on the numerical algorithms that can be used in the discrete setting to compute these atoms.)

Let us now consider the following synthetic image:

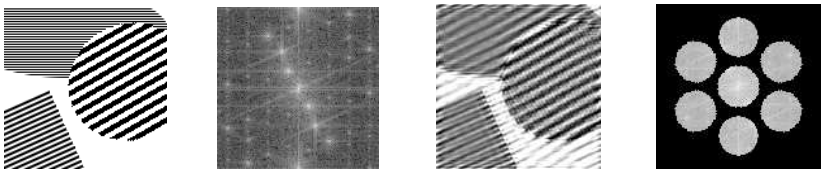

Fig. 7. From left to right: the original 128×128 image g_0 and its spectrum $\log|\mathcal{F}(g_0)|$, the corrupted image g and its spectrum $\log|\mathcal{F}(g)|$.

In Fig. 8, we can see that these atoms behave as Gabor atoms by capturing different patterns in the corrupted image. The different regions of g can thus be distinguished by analyzing the filtered $g * \phi_n$.

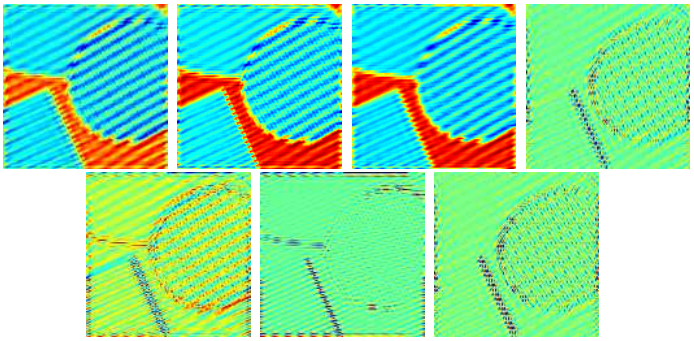

Fig. 8. Filtered $g * \phi_n$, $n = 1, ..., 7$.

The previous atoms were computed for the mask considered in Fig. 4. However, our approach is quite general and can be adapted to any other corruption

mask. As an example, here follow the atoms we get if we consider the mask that comes into play in the scattering problem:

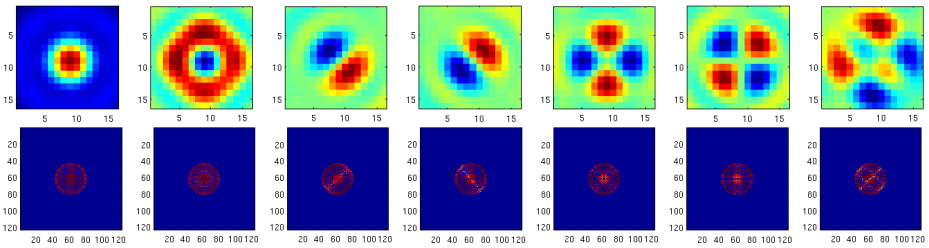

Fig. 9. First line: Zoomed-in atoms $(\phi_n)_{n=1,\dots,7}$ with $p = 4$ adapted to the scattering problem. Second line: Their respective spectra.

These atoms are also really localized and are about the same size as a classical image patch (about 10 pixels wide) which therefore allows us to correlate local features of an image within a Non-Local approach.

6 Distance Map Comparison

We are going to compare the performance of the different similarity measures considered so far by fixing one patch of the *corrupted image* g (indicated in green in the clean image) and identify the 13 best matches in the *corrupted image* (and indicated in red in the clean image) according to

1. The atom-based distance $\delta^1(x_k, x_\ell) = \left(\sum_{n=1}^{7} |g * \phi_n(x_k) - g * \phi_n(x_\ell)|^2\right)^{\frac{1}{2}}$.
2. The ideal distance $\delta^2(x_k, x_\ell) = \|p_{g_0}(x_k) - p_{g_0}(x_\ell)\|_2$ based on the clean g_0.
3. The SSD $\delta^3(x_k, x_\ell) = \|p_g(x_k) - p_g(x_\ell)\|_2$.

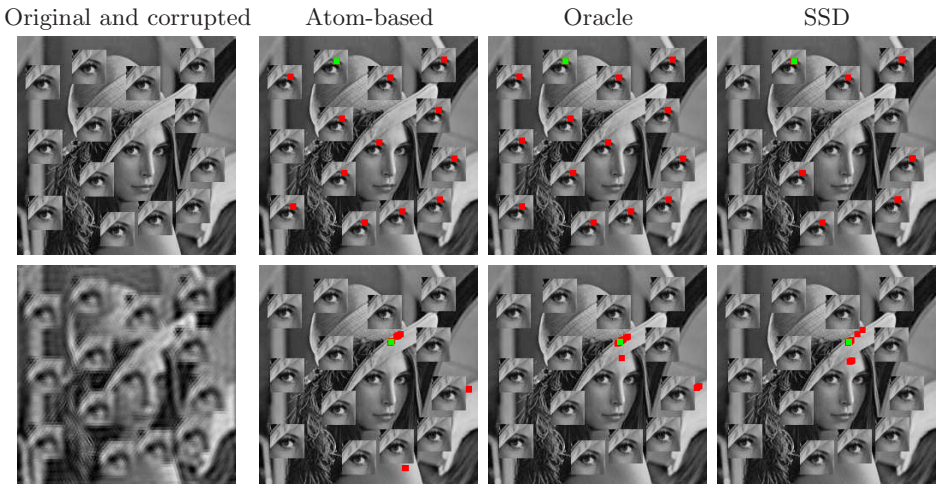

Fig. 10. Best matches (in red) corresponding to a fixed patch (depicted in green) in an image with lot of self-similarities.

The two lines correspond to two different choices for the patch or reference (first line: a small square on the eyebrow from the second large eye patch, second line: the edge of the hat towards the middle) The second and fourth columns are computed using the corrupted image g. The results are displayed on the clean image. In the third column we depicted the results obtained with the SSD computed with the clean image.

These two examples suggest that the atom-based distance we designed performs better than the SSD in some cases: indeed, the atom based results are closer to the "oracle" than the SSD which looses more matches.

7 Numerical Experiment

From now on, we are going to make important assumptions that let us drastically improve the complexity of the minimization problem (11) and save memory:

(*i*) Two patches are unlikely to be the same if they are far from each other thus we can assume that for any fixed x_k, x_ℓ is a candidate if $|x_k - x_\ell| \leq \eta$ for some fixed η. This defines a neighborhood of candidates (also called *window*).

(*ii*) Given x_k, x_ℓ, minimizing (10) yields a $v_{k,\ell}$ whose support is actually $(x_k + \mathrm{supp}(\psi)) \cup (x_\ell + \mathrm{supp}(\psi))$. Therefore, assuming that $|x_k - x_\ell| \geq \varepsilon$ for some $\varepsilon \leq \rho$, we can still get a global minimizer for (11). Although, usually $\varepsilon = 1$ pixel, assuming that $\varepsilon = 7$ pixels for a patch size $\rho = 7$ or 9 pixels (to let the reconstructed patches to overlap) results in an acceleration of order 10.

(*iii*) Once w computed, we can for any fixed pixel x_k keep only the m_0 best matches x_ℓ (see Fig. 10).

Weight computation: Notice that for a fixed mask M, the atoms $(\phi_n)_{n=1,\dots,n_0}$ can be computed in advance and stored. Then $g_n = \phi_n * g$ can be readily, computed once for all, using the `fft`. Now, we recall that

$$\left(\delta^1(x_k, x_l)\right)^2 = \sum_{n=1}^{n_0} |g_n(x_k) - g_n(x_\ell)|^2, \tag{16}$$

$$\left(\delta^3(x_k, x_l)\right)^2 = \sum_{i=1}^{\rho^2} |\left(p_g(x_k)\right)_i - \left(p_g(x_\ell)\right)_i|^2. \tag{17}$$

From the previous definitions, it is readily seen that if $\rho = 7$ or 9 pixels is the patch size, then, as far as $(n_0 \leq \rho^2)$ [1], it is faster to compute δ^1. In practice, this can result in an acceleration of order 10. Such a gain in complexity will be observed numerically in the sequel.

Algorithmic issues: An interesting feature of our work is that most optimization problems which are solved are constrained quadratic minimizations. Thus,

[1] which will always be the case, otherwise it means that all atoms are essentially compactly supported which is possible only if the mask is full

the best choice (which we used) is the conjugate gradient. We also used Beck-Teboulles's algorithm [2] in some instances, since our same program was optionally implementing an additional TV-regularization, with no real difference. The assumptions $(i) - (iii)$ allow us to accelerate the computation.

Numerical results: In the following tests we consider three different energies:

- The constrained total variation minimization problem

$$\min_{v \in \mathcal{M}^\perp} TV(g+v),$$

  which can be solved for instance using [2, 25] or even [4].
- The constrained minimization problem (11) where we consider the SSD measure of similarity δ^3. In the simulations of Fig. 12 and Fig. 13, we take the window size $\eta = 20$ pixels, the patch size $\rho = 7$ pixels and $\varepsilon = 5$ pixels. For a fixed patch, we only keep the $m_0 = 10$ best matches and $h = 100$. For the tests of Fig. 14-18, we set $\eta = 100$, $\rho = 5$, $\varepsilon = 3$, $m_0 = 6$, $h = 100$. For the smaller example of Fig. 11 we chose $\eta = 20$, $\rho = 5$, $\varepsilon = 1$, $m_0 = 8$, $h = 100$. In Fig. 19 we chose $\eta = 60$, $\rho = 9$, $\varepsilon = 3$, $m_0 = 10$, $h = 100$.
- The constrained minimization problem (11) where the weight is computed with the atom-based distance δ^1. In the tests of Fig. 12 and Fig. 13, we consider the $n_0 = 25$ first atoms with a moment $p = 4$. As above we take $\varepsilon = 5$ pixels, $h = 100$ and for a fixed patch, we only consider the $m_0 = 10$ best matches. For Fig. 14-17, we set $\eta = 100$, $\varepsilon = 3$, $m_0 = 6$, $n_0 = 18$, $h = 100$, $p = 4$. To produce Fig. 19 we took $\eta = 60$, $\varepsilon = 3$, $m_0 = 10$, $n_0 = 18$, $h = 100$, $p = 4$. For the small image of Fig. 11 we picked $\eta = 20$, $\varepsilon = 1$, $m_0 = 8$, $n_0 = 18$, $h = 100$, $p = 20$.

In our first example (Fig. 11), we consider a *small* 64×64 image and both exponents $\alpha = 1$ and $\alpha = 2$ in equation (11):

For this first example, both choices for the parameter α provide improved results. The minimization is possible either by approximating the ℓ^1 norm or by using modern splitting algorithms as introduced in [28, 8].

In this experiment, it is interesting to observe that the robust choice $\alpha = 1$ is particularly useful with the SSD distance which is likely to produce outliers (spurious matches), and almost useless with the oracle distance (see Fig. 5) which produces perfect matches. Our distance seems to lie in-between, in a region where the choice is apparently less decisive.

Henceforth, we shall only consider the exponent $\alpha = 2$ which is numerically tractable for larger images (the case $\alpha = 1$ being in practice, as expected, quite slower than the quadratic case). The experiment that follows (Fig. 12) is a simple but larger 128×128 toy example.

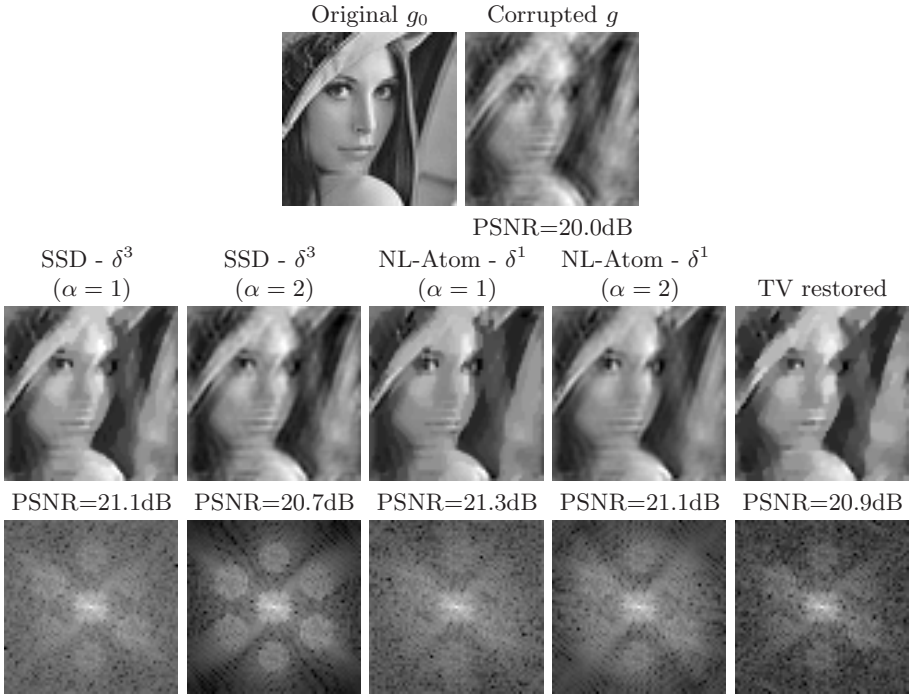

Fig. 11. Corrupted and restored 64×64 *Lena* and their respective spectra for exponent values $\alpha \in \{1, 2\}$.

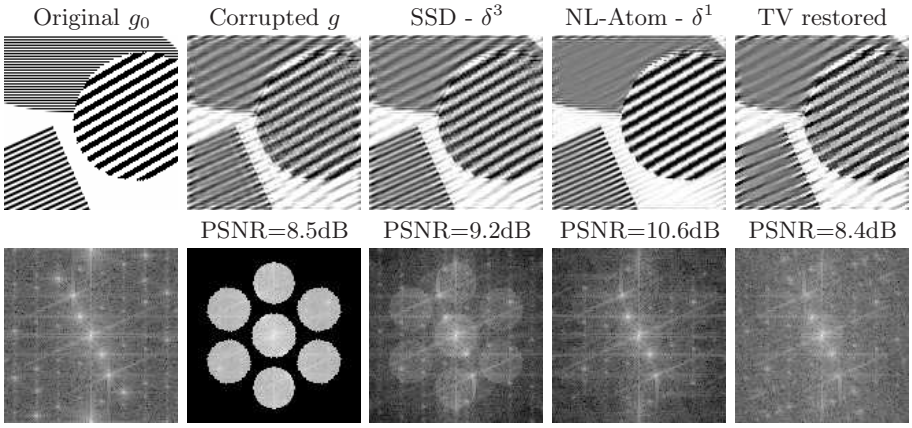

Fig. 12. Corrupted and restored 128×128 toy example and their respective spectra.

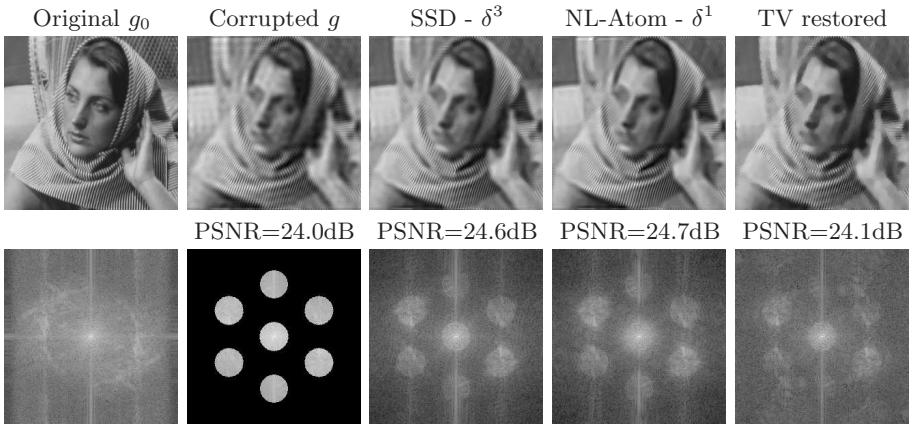

Fig. 13. Corrupted and restored 256×256 crop of *Barbara* and their respective spectra.

Fig. 12 and 13 show an example with the Mask of Fig. 4. In this case, the best results are obtained with the atom-based distance. In particular in Fig. 12, the oscillations with frequencies in the mask are almost perfectly recovered and the spurious contamination is very low.

Let us see how these methods perform for the acoustic scattering problem of Section 1. To do so, we consider the mask introduced in Fig. 1.

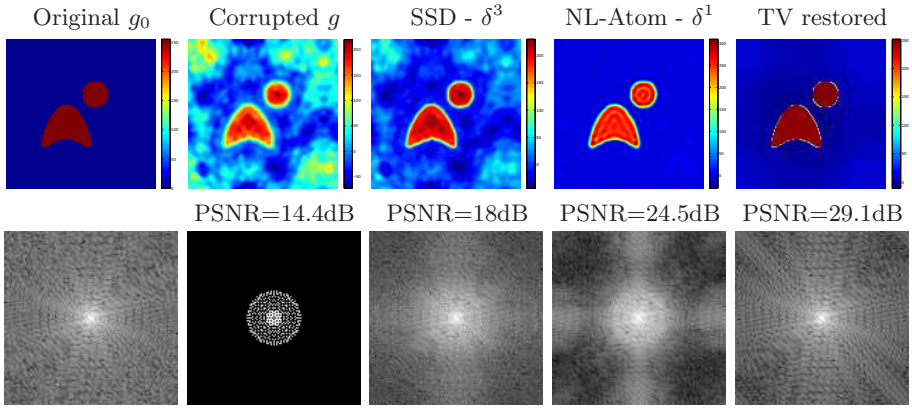

Fig. 14. Corrupted and restored scatterers and their respective spectra.

In this case the best results are obtained with the standard Total Variation: this is consistent with the fact that the object to recover is piecewise constant and the loss is mostly in the high frequencies (for further details on these qualitative properties see [15–17]).

In the next simulation, we are going to assume that the Fourier coefficients that we kept are contaminated by a Gaussian noise of magnitude $0.03\|g_0\|_2$.

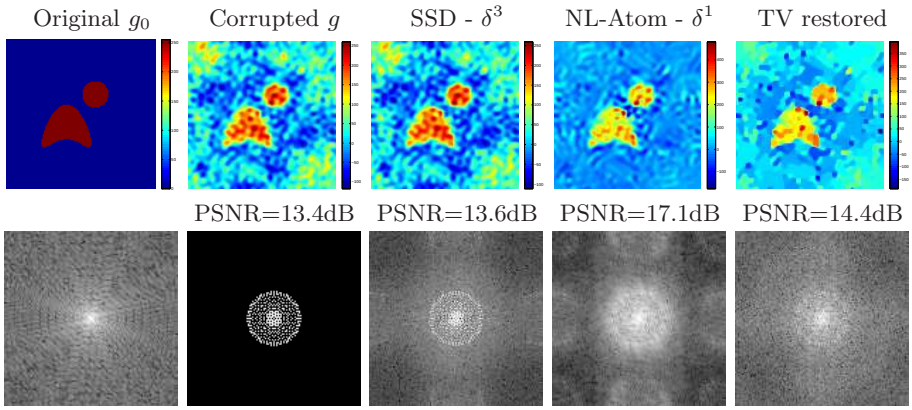

Fig. 15. Corrupted and restored scatterers and their respective spectra.

In this case, the Total Variation based optimization find it much harder to recover the data and the Atom-based distance yields satisfactory results. Now considering the Born approximation (2), let us use the data that comes out of the direct problem. In a sense, this amounts to adding to the Fourier coefficients a noise whose distribution is unknown. Our method is able to get rid of all the spurious objects.

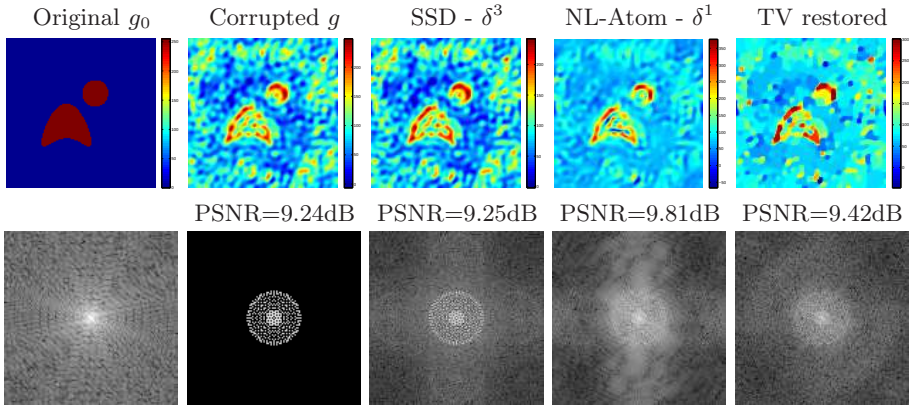

Fig. 16. Corrupted and restored scatterers and their respective spectra.

In this kind of problems it is usually important to distinguish objects that are very close. Let us consider such a situation:

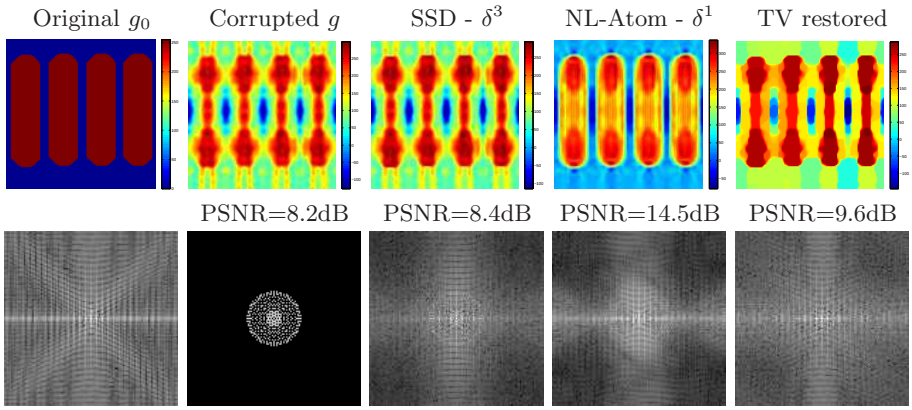

Fig. 17. The original image 128×128 is made of 4 scatterers separated by 6 pixels. We then depicted respectively the corrupted and the restored scatterers and their spectra.

Actually, in the latter experiment, we were able to distinguish objects that are separated by 0.56λ which is quite close to the theoretical limit 0.5λ. We were able to achieve this result since, we recall, two identical patches in the original image g_0 do not necessarily match according to the SSD but do match according to the atom-based distance. This explains why in Fig. 17, the atom-based distance improves the matching step (as in Fig. 10) and provides logically better results as observed in the different tests. In this case, there is to our knowledge no method but the one we propose to improve the results over the raw data.

Weight recomputation: In [11, 12, 26, 27, 1], it was pointed out that one can get improved results (especially for inpainting problems) by allowing recomputation of the SSD-based weight on the being restored image. We depict in Fig. 18 the restorations one can expect after several recomputations.

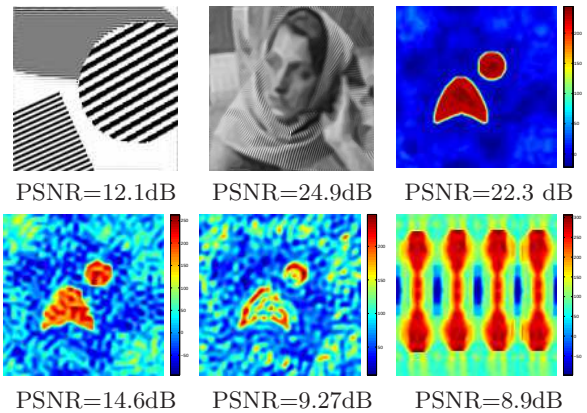

Fig. 18. Restored images after many SSD weight recomputations.

This procedure is really cumbersome and does not always improve results over the atom-based method. The weight recomputation is not possible for the atom-based distance we introduced since the distance computed on the restored image is exactly the same as the one computed on the corrupted image. However, our method can be used as an initialization for the classical weight recomputation to improve results further. This is the strategy we adopt in the following tomography problem where the Fourier coefficients got corrupted by a Gaussian noise of magnitude $0.3\|g_0\|_2$:

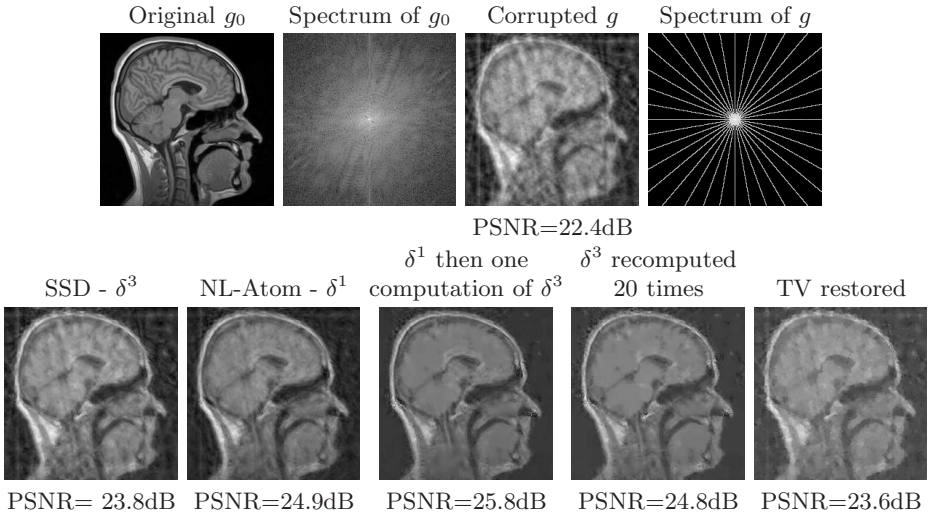

Fig. 19. Restoration for a 240×240 tomography image.

Computation Times: As we claimed previously in our discussion on complexity, the computation time with our approach (which in a first step reduces the dimensionality to retain only the information which is useful for the matching) is quite reduced. For instance, in Fig. 19 we observed the following computation times

Approach	Computation time
SSD - δ^3	83s (dist) + 15s (iter) = 98s
NL-Atom - δ^1	**8s (dist) + 25s (iter) = 33s**
δ^1 then one δ^3	143s
δ^3 recomputed 20 times	>30min
TV restored	82s

We performed our experiments in Matlab, so the absolute CPU times are not really relevant.

8 Conclusion

In this paper, we have considered the problem of reconstructing an image with a known pattern of missing Fourier coefficients, by means of a non-local method

which assumes and exploits some spatial redundancy of the original image. In order to detect similar patches on the original image, we have introduced an original similarity criterion which is different from a standard quadratic distance, and is insensitive to the image degradation. This distance is based on precomputed atoms which are then used to filter out the degradation of the image, without destroying the local features.

Then, by minimizing a simple variational model (based on a quadratic energy, which we claim would produce the exact solution if the similar patches were exactly known—and in large enough quantity), we have experimentally shown the efficiency of this non-local framework, for some Fourier pattern. If the loss of coefficients is only in the high frequencies, then our approach does not produce any improvement over standard zooming techniques (local or nonlocal). On the other hand, for more complex masks which miss low frequency Fourier coefficients, our results are far superior to results obtained with more usual measures of redundancy.

References

1. P. Arias, G. Facciolo, V. Caselles, and G. Sapiro, *A variational framework for exemplar-based image inpainting*, International journal of computer vision, (2011).
2. Amir Beck and Marc Teboulle, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM J. Imaging Sci., 2 (2009), pp. 183–202.
3. A. Buades, B. Coll, and Jean-Michel Morel, *A non-local algorithm for image denoising*, in CVPR 2005, vol. 2, IEEE, 2005.
4. Antonin Chambolle and Thomas Pock, *A first-order primal-dual algorithm for convex problems with applications to imaging*, J. Math. Imaging Vision, 40 (2011), pp. 120–145.
5. Tony Chan, Jianhong Shen, and Hao-Min Zhou, *Total variation wavelet inpainting*, J. Math. Imaging Vision, 25 (2006), pp. 107–125.
6. T. Chan and L.A. Vese, *Active contours without edges*, Image Processing, IEEE Transactions on, 10 (2001), pp. 266–277.
7. David Colton, *Inverse acoustic and electromagnetic scattering theory*, in Inside out: inverse problems and applications, vol. 47 of Math. Sci. Res. Inst. Publ., Cambridge Univ. Press, Cambridge, 2003, pp. 67–110.
8. Patrick L Combettes and Jean-Christophe Pesquet, *Proximal splitting methods in signal processing*, in Fixed-Point Algorithms for Inverse Problems in Science and Engineering, Springer, 2011, pp. 185–212.
9. Kostadin Dabov, Alessandro Foi, Vladimir Katkovnik, Karen Egiazarian, et al., *BM3D image denoising with shape-adaptive principal component analysis*, in SPARS'09-Signal Processing with Adaptive Sparse Structured Representations, 2009.
10. AJ Devaney, *Inversion formula for inverse scattering within the born approximation*, Optics Letters, 7 (1982), pp. 111–112.
11. G. Gilboa, J. Darbon, S. Osher, and T. Chan, *Nonlocal convex functionals for image regularization*, UCLA CAM Report, (2006), pp. 06–57.
12. Guy Gilboa and Stanley Osher, *Nonlocal operators with applications to image processing*, Multiscale Model. Simul., 7 (2008), pp. 1005–1028.

13. G. Gilboa, N. Sochen, and Y.Y. Zeevi, *Variational denoising of partly textured images by spatially varying constraints*, IEEE Transactions, 15 (2006), pp. 2281–2289.
14. F. Guichard and F. Malgouyres, *Total variation based interpolation*, in Proceedings of the European signal processing conference, vol. 3, 1998, pp. 1741–1744.
15. Khalid Jalalzai, *Discontinuities of the minimizers of the weighted or anisotropic total variation for image reconstruction*, (2012).
16. ———, *Some remarks on the staircasing phenomenon in total-variation based image denoising*, (2012).
17. ———, *Regularization of inverse problems in image processing*, PhD Thesis, École Polytechnique, Palaiseau, (March 2012).
18. A.M. Johnson, RJ Eastwood, and AH Greenaway, *Optical aperture synthesis*, in 3rd Electro Magnetic Remote Sensing Defence Technology Centre (EMRS DTC) Technical Conference, Edinburgh, Scotland, July, 2006, pp. 13–14.
19. C. Kervrann and J. Boulanger, *Optimal spatial adaptation for patch-based image denoising*, Image Processing, IEEE Transactions on, 15 (2006), pp. 2866–2878.
20. Stefan Kindermann, Stanley Osher, and Peter W Jones, *Deblurring and denoising of images by nonlocal functionals*, Multiscale Modeling & Simulation, 4 (2005), pp. 1091–1115.
21. A. Kirsch, *Characterization of the shape of a scattering obstacle using the spectral data of the far field operator*, Inverse problems, 14 (1998), p. 1489.
22. Stéphane Mallat, *A wavelet tour of signal processing*, Elsevier/Academic Press, Amsterdam, third ed., 2009. The sparse way, With contributions from Gabriel Peyré.
23. Y. Meyer, *Du nouveau sur l'échantillonnage ou comment tricher avec shanon*, (2009).
24. Lionel Moisan, *Extrapolation de spectre et variation totale pondérée*, in 18ème Colloque sur le traitement du signal et des images, FRA, 2001, GRETSI, Groupe d'Etudes du Traitement du Signal et des Images, 2001.
25. Yurii Nesterov, *Gradient methods for minimizing composite objective function*, ReCALL, 76 (2007).
26. G. Peyré, S. Bougleux, and L. Cohen, *Non-local regularization of inverse problems*, Computer Vision–ECCV 2008, (2008), pp. 57–68.
27. Gabriel Peyré, Sébastien Bougleux, Laurent D Cohen, et al., *Non-local regularization of inverse problems*, Inverse Problems and Imaging, 5 (2011), pp. 511–530.
28. Hugo Raguet, Jalal Fadili, and Gabriel Peyré, *Generalized forward-backward splitting*, arXiv preprint arXiv:1108.4404, (2011).
29. P.M. Van den Berg and RE Kleinman, *A total variation enhanced modified gradient algorithm for profile reconstruction*, Inverse Problems, 11 (1995), p. L5.
30. X. Zhang and T.F. Chan, *Wavelet inpainting by nonlocal total variation*, CAM Report (09-64), (2009).